DOCUMENT RESUME

ED 296 587 FL 017 468

AUTHOR Read, Walter; And Others
TITLE What's So Hard about Understanding Language?
PUB DATE Apr 88
NOTE 7p.; Paper presented at the Annual Meeting of the American Educational Research Association (New Orleans, LA, April 5-9, 1988).
PUB TYPE Information Analyses (070) -- Speeches/Conference Papers (150)

EDRS PRICE MF01/PC01 Plus Postage.
DESCRIPTORS *Ambiguity; *Artificial Intelligence; Comprehension; *Epistemology; *Inferences; *Language Processing; Logical Thinking; Models; Semantics
IDENTIFIERS *Anaphora; Metonymy; *Natural Language

ABSTRACT

A discussion of the application of artificial intelligence to natural language processing looks at several problems in language comprehension, involving semantic ambiguity, anaphoric reference, and metonymy. Examples of these problems are cited, and the importance of the computational approach in analyzing them is explained. The approach applies specific mechanisms or models to natural language to make explicit the complex system of assumptions, knowledge, and processes used in understanding language. It is proposed that by examining language at the micro level it is possible to see how much a learner--whether a computer, a child, or a second language learner--must master. (MSE)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

ED296587

What's So Hard About Understanding Language?

Walter Read
Michael Dyer
Artificial Intelligence Laboratory
UCLA Computer Science Department

Richard Feifer
Center for Technology Assessment
UCLA Graduate School of Education

"PERMISSION TO REPRODUCE THIS MATERIAL HAS BEEN GRANTED BY
W. Read
TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)."

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)
This document has been reproduced as received from the person or organization originating it.
Minor changes have been made to improve reproduction quality.
Points of view or opinions stated in this document do not necessarily represent official OERI position or policy.

Paper presented at the Annual Meeting of the
American Educational Research Association

for the Symposium
"Understanding Natural Language Understanding:
Helping Computers to Think Like People"

April, 1988, New Orleans

(Please do not quote or cite without prior written consent of author)

FL017468

Overview

This paper reports empirical work in artificial inteligence (AI) on natural language processing (NLP). AI has come to mean many things to many people. On one extreme we have engineers who are primarily interested in software tools with which to build computer programs. On another extreme, AI touches on issues in philosophy of mind and philosophy of science that often rouse strong emotions.

The kind of AI that we are interested in is cognitive modelling. The goal here is to describe some aspect of human mental activity, say problem solving or language use, in computational terms. The central issues is this kind of study turn out to be representation and reasoning (corresponding roughly to data structures and control in computer programs). By representation we mean the content and organization of information used. By reasoning we mean not only inference in the traditional sense but generally any method for drawing conclusions from the data. This includes analogical reasoning, heuristics for educated guesses, methods for deciding between conflicting evidence, etc. The details of representation and reasoning are a major concern of AI workers and controversies rage among proponents of different schools. Some researchers argue that only when the representation is given in predicate logic and the reasoning is reducible to inference in some formal sense can an AI model be considered seriously. Others argue that the representation should be given in a network which resembles a neural network and that reasoning should be activation (with, perhaps, some inhibition) over the network. For our present purposes, these implementation details are not important. What is important is what knowledge is needed and what is done with it.

The AI-NLP Enterprise

In this paper we will examine closely several problems in language understanding involving semantic ambiguity, anaphoric reference, and metonymy. These examples come from our own work and the work of other AI researchers.

Ambiguity

Consider the following examples:

1(a). John smoked the pot.

1(b). John put the pot in the dishwasher.

1(c). John heard the police coming. He quickly hid the pot in the dishwasher.

The issue here is disambiguating *pot* as 'cooking pot' or as 'marijuana'. Assuming that the listener has disambiguated *smoked* (as 'burned in order to inhale' rather than 'prepared food by exposing to smoke' or other more informal uses), the first example seems straightforward. Only the 'marijuana' sense satisfies the selectional restrictions on the appropriate meaning of 'smoked'. The second example is more difficult. Either meaning of 'pot' fits the selectional restrictions and, of course, the sentence is ambiguous. Nonetheless, when this sentence is heard in isolation, most people prefer the 'cooking pot' meaning. Presumably people have knowledge of the normal uses of a dishwasher. The third example is the most complex. It depends on knowing the usual role of police and the social and legal status of marijuana versus cooking pots. In a society where marijuana was legal but pots were valuable (and therefore possibly stolen) we might disambiguate 'pot' differently. Examples like this raise difficult issues of meaning. The fact that marijuana is currently illegal in this society isn't clearly a "compositional" feature of its meaning. In this story, 'dishwasher' is also ambiguous (between person and machine) and is disambiguated by 'hid ... in'. Understanding this story then depends on a less-obvious aspect of dishwashing machines, namely that they may be treated as opaque containers.

The point here is that ambiguity is common and even in simple cases requires quite a bit of world knowledge to resolve. In building systems that can handle this kind of ambiguity, we quickly encounter problems in organizing the necessary knowledge and in activating the "right" pieces. Why is the *illegality* of marijuana relevant for the third example and not for the first two? Can a dishwasher *only* be used on cooking pots and not on marijuana? How do we know that? (I.e., what kind of organization of knowledge or reasoning allows us to make the right conclusion?) The reasoning needed to choose the appropriate evidence and to choose a preferred meaning is complex.

Metonymy

Metonymy is in general the use of one thing to refer to something related, as in saying 'the sword' to mean military action or in referring to a diner customer as 'the ham sandwich'. It is often treated as a literary device but in fact is in common use even in technical contexts. In their study of ill-formed input, Carbonell and Hayes (1983) examine a front end to a database of computer equipment in which the user says:

2. Can you connect a video disk drive to the two megabytes?

Of course, the author doesn't want to add the drive to the count or even the memory being counted. The drive is to be added to the computer (mentioned earlier in the discourse) that has the two megabytes (as compared to some other computer mentioned). These examples are easy for people knowledgible in the area to understand but it is difficult to formulate any general algorithm describing how we do it. People without specific knowledge of computers might find the discourse incomprehensible.

Anaphoric reference

Not only do we need appropriate representations to process the sentences we hear, but we also need to be able to modify the representations in the course of listening. Consider the following situations:

3(a). On the shelf, I have a hat from the Super Bowl and a cup from Chicago. These things are my newest souvenirs.

3(b). I found a nickel, a dime and a quarter in the drawer. The oldest of these was from 1919.

In the first example, a listener must resolve the reference of 'these things'. A simple heuristic is to track back to the most recent *set* of possible referents that can be considered 'things'. But no such set has been mentioned. An understander (whether natural or artificial) must create the set from an appropriate list of things. In the second example, 'oldest' clearly refers to a set that is ordered by time. The understander must not only form the set but must know enough about the objects to realize that a time order is possible.

It is not obvious what kind of a strategy gives us the right representation. Given the enormous number of possible sets that we might have to create and the enormous number of ways to construe them, it seems unlikely that we store every possible combination of the things we know. So we must be able to create at least some of the knowledge structures in response to the conversational demands. But what process allows us to build the right sets? In example 3(a), the things are the hat and the cup, but not the shelf, so it can't be that we just form groups of things. The language processor should know something about souvenirs and the kinds of things that get chosen. But the shelf might be a souvenir from a time of interest in woodworking. In example 3(b), the age of the drawer isn't important, is it? A language understander also must understand when an object (or action or set of objects or ...) is in *focus*. Focus is a function of many things, including our knowledge of how language is used in interaction and knowledge about the social roles of objects and actions. Understanding focus seems to be to a great extent cultural.

The Importance of the Computational Approach

Much of the work we do in AI looks simple. Certainly, none of these sentences seem particularly difficult to understand. But in trying to build an actual artificial understander we are forced to make explicit the assumptions, knowledge and processes we use. As these examples show, even simple sentences can require a lot of world knowledge and non-trivial reasoning strategies. Of course, we are interested in finding general rules for language understanders. What form of organization of the knowledge allows us to access the information we need when we need it? What heuristics do we use to choose the most useful features of the knowledge? What assumptions do we need to make about other speakers and listeners in order to communicate?

This work is interesting because by insisting on specific mechanisms we can make explicit the enormous amount of work we do thoughtlessly when we engage in even simple reading or conversation. Because our years of experience in thinking, reasoning, persuading, etc. allow us to process language almost automatically it is easy to forget just how complicated this processing is. It is easy and tempting to try to reduce this complexity by assuming a few simple, formal rules or by assuming that what is obvious to us must be obvious to everybody else. It is easy to underestimate the sophistication of the learning and reasoning strategies of small children, or even infants.

This kind of thinking is possibly more dehumanizing than the technology we often worry so much about. By examining cognitive processing at the micro level we can see how much a learner, whether a computer, a child or a second language learner, has to master.

Beyond the details of language understanding, the surest lesson from this kind of research is that each one of us is extremely complex and capable of ordinary feats that still stump the best computers.

Reference

Carbonell, James G. & Hayes, Philip J. (1983). Recovery strategies for parsing extragrammatical language. *American Journal of Computational Linguistics, 9*, 123-146.